

# Riemannian Geometry Speaks Louder Than Words: From Graph Foundation Model to Next-Generation Graph Intelligence

Philip S. Yu

University of Illinois Chicago  
Chicago, USA  
psyu@uic.edu

Li Sun

Beijing University of Posts and Telecommunications  
Beijing, China  
lsun@bupt.edu.cn

## Abstract

Graphs provide a natural description of the complex relationships among objects, and play a pivotal role in communications, transportation, social computing, the life sciences, etc. Currently, there is strong agreement that Graph Foundation Models (GFMs) are essential for advancing graph learning, yet considerable disagreement persists on how to build a powerful, general-purpose GFM analogous to Large Language Models (LLMs). Graph Neural Networks (GNNs) exhibit limitations in memory retention and principled interpretability when confronted with multi-domain pretraining and adaptation. The challenge of graph serialization hinders the direct application of LLMs, as the words struggle to capture the structural complexity and diversity inherent in graphs. In contrast, Riemannian geometry offers an elegant mathematical framework for modeling structures, while remaining compatible with graph semantic learning, even with LLMs. In this paper, we argue that, for graphs, **Riemannian geometry speaks louder than words**, and lay out the foundational principles for GFM. Reimagining with Riemannian geometry, we introduce a blue sky idea—**Riemannian Foundation Model (RFM)**—that opens a new pathway for capturing complex structural patterns and uncovering cross-domain generalities. RFM emphasizes intrinsic graph geometry and embodies endogenous capacities for structural inference and generation, moving beyond mere representation-space switching. Accordingly, we outline a progressive agenda that begins with universal structural understanding through intrinsic geometry, and then rebuilds LLM with a Riemannian engine for general-purpose graph modeling and beyond. Thus, RFM enables a paradigm shift from designing graph models to solving graph-structured applications with RFM agents, unlocking the next-generation graph intelligence.

## CCS Concepts

• **Information systems** → **Data mining**; • **Computing methodologies** → **Neural networks**; *Unsupervised learning*; • **Mathematics of computing** → *Topology*.

## Keywords

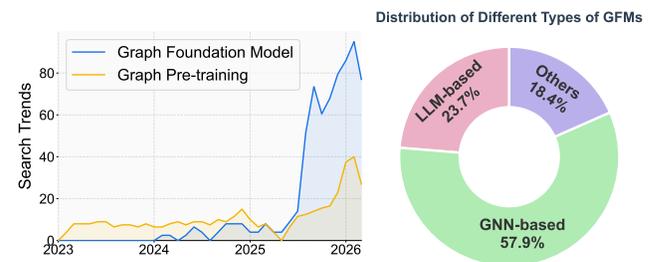
Graphs, Topologies, Riemannian Geometry, Foundation Models

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
*SIGKDD'26, Jeju, Korea*

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
ACM ISBN 978-1-4503-XXXX-X/2018/06  
<https://doi.org/XXXXXXX.XXXXXXX>

## ACM Reference Format:

Philip S. Yu and Li Sun. 2026. Riemannian Geometry Speaks Louder Than Words: From Graph Foundation Model to Next-Generation Graph Intelligence. In *Proceedings of 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining (SIGKDD'26)*. ACM, New York, NY, USA, 7 pages. <https://doi.org/XXXXXXX.XXXXXXX>



**Figure 1: Left: Google Trends for Graph Foundation Model and Graph Pre-training. Right: Proportion of existing graph foundation model papers by architecture type.**

## 1 Introduction

Graphs are ubiquitous, and graph learning is pivotal to the advancement of data mining and artificial intelligence. The emergence of Large Language Models has catalyzed a profound paradigm shift in this domain, driving the field from specialized, fragmented models toward universal foundation models. Recent efforts to replicate the success of foundation models in graph learning have ignited growing interest in Graph Foundation Models—a class of expressive, general-purpose, pretrainable neural architectures for graphs, as illustrated in Figure 1.

Unlike texts and images, the non-Euclidean nature of graphs charts a distinct developmental course. The road to GFM is currently divided as in Figure 1, with both prevailing approaches facing inherent challenges. One line of work extends Graph Neural Networks via multi-domain pretraining [23, 53]. However, these models are fundamentally constrained by local message passing and fixed inductive biases, such as Euclidean representation spaces, failing to capture structural diversity. Their limitation in memory retention also impedes the discovery of cross-domain generalities. Another line of work leverages LLMs as the cornerstone, demonstrating initial success primarily on text-attributed graphs [27]. This paradigm focuses on devising word-like graph serialization techniques to accommodate LLMs, yet it inherently struggles to traverse the vast structural diversity inherent in graphs [58, 60]. Also, representing graphs as “word” sequences poses a fundamental challenge in capturing their intricate structural complexity [67].

We argue that, **for graphs, Riemannian geometry speaks louder than words**. To date, neither LLM-based nor GNN-based

**Table 1: Capability Matrix of Graph Foundation Models (GFMs) across Diverse Tasks**

	<i>OpenGraph</i> [52]	<i>OFA</i> [27]	<i>GraphAry</i> [64]	<i>RAGraph</i> [18]	<i>GCOPE</i> [63]	<i>GRAVER</i> [58]	<i>GIT</i> [49]	<i>RiemannGFM</i> [35]	<i>BRIDGE</i> [59]	<i>GraphGPT</i> [40]	<i>AnyGraph</i> [51]	<i>AutoGFM</i> [5]	<i>SA<sup>2</sup>GFM</i> [32]	<i>MDGPT</i> [57]	<i>SAMGPT</i> [56]	<i>FedGFM+</i> [68]	<i>FedBook</i> [50]	<i>UniGraph</i> [14]	<i>MDGCL</i> [66]	<i>GFT</i> [48]
<i>Node Classification</i>	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
<i>Link Prediction</i>	✓	✓					✓	✓		✓	✓					✓	✓	✓		✓
<i>Graph-level Task</i>	✓	✓	✓	✓		✓	✓	✓	✓			✓	✓	✓	✓	✓	✓	✓	✓	✓
<i>Graph Generation</i>																				

approaches have identified the graph generalities or invariances that underpin universal graph modeling. Standing at a crossroads, we return to the essence of graphs—the discrete analogs of manifolds—and thereby illuminate a new direction: a Riemannian geometry approach. Indeed, Riemannian geometry provides a rigorous foundation for graph modeling. In this elegant framework, there exists a powerful toolkit to build GFMs, including: 1) a rich family of manifolds to capture structural diversity, 2) the concept of curvature to delve into structural complexity, 3) the moving frame method [33] to transfer knowledge across domains, and 4) the construction of vector bundles [43] to integrate graph structure and semantics, creating a synergistic interface to integration with LLMs.

Grounded in Riemannian geometry, we put forward a blue—sky idea of **Riemannian Foundation Model**, aimed at general-purpose graph learning and beyond. While a significant body of work on non-Euclidean graph learning exists, it typically focuses on selecting extrinsic representation spaces, such as hyperbolic spaces, tailored to specific graphs [34, 36, 45]. Moving beyond the mere switching of representation spaces, RFM emphasizes graph intrinsic geometry such as flatness, symmetry, connectivity, and anisotropy. It offers a principled framework to uncover graph generalities, interpret knowledge transfer across domains, and endow models with endogenous capacities for structural inference and generation. To realize this vision, we outline a progressive agenda toward RFM, thereby opening the next-generation graph intelligence. Analogous to agents for LLMs, RFM agents are envisioned as the key to solving graph-structured applications, from code generation to drug design.

## 2 Current Challenges

### 2.1 Are GNNs Generalizable?

Graph Neural Networks, powered by the message-passing mechanism, have become the dominant solution for graph representation learning [6, 13, 21, 44]. Recently, a surge of research has sought to extend GNNs into the realm of foundation models [27, 27, 52, 63], yet a fundamental mismatch persists between the two. While GNNs are typically trained on specific graphs by task-specific objectives, Graph Foundation Models aspire to achieve general-purpose modeling across arbitrary graphs through large-scale pretraining. With the advent of self-supervised techniques such as graph contrastive learning [1, 16, 49], GNNs have reduced their dependence on target tasks by learning from the data themselves. Nevertheless, GNNs

still struggle with structural shifts in data and suffer from inherent scalability issues, e.g., the well-known oversmoothing issue preventing the network from going deep [19].

We posit that the core challenge lies in accommodating the vast diversity of the graph domain. Graphs exhibit wildly different topological patterns, ranging from the rigid bonds in molecules to the chaotic links in social networks [22, 39, 61], yet conventional GNNs rely on a single, fixed Euclidean representation space. This severely limits their expressiveness in differentiating structural patterns. The problem is exacerbated during cross-domain generalization. While multi-domain pretraining has been proposed to enhance generalization [15, 46, 56, 59], catastrophic forgetting remains problematic. Notably, jointly capturing shared knowledge while preserving distinct structural semantics remains inherently challenging [35], and a principled interpretive framework is still lacking.

We argue that, **although multi-domain pretraining marks progress, the GNN mechanism itself reflects a deeper inability to represent diverse topologies and the structural semantics they encode.**

### 2.2 Can LLMs Capture Graph Geometry?

LLMs have revolutionized textual understanding and generation, prompting a paradigm shift toward developing LLM-based GFMs. A central question arises: can linguistic tokens—“words”—adequately capture the structural complexity inherent in graphs? Within this line of research, much effort has focused on graph serialization to ensure compatibility with LLM architectures. Existing serialization methods include graph random walks [17, 37, 40, 55, 62], path sampling [20, 42], and Hamiltonian cycles rooted in topological theory [3]. These methods aim to equip LLMs with structural knowledge by exposing them to local graph contexts. However, a traversal offers only a fragmented, path-dependent view of the graph, while falls short of capturing the high-order structural patterns that govern the entire topology. We argue that **universal graph modeling via LLMs remains a distant goal, hindered by the intractability of traversing diverse graph domains and the limitation of representing structural semantics through words alone.**

Mimicking the advances in LLMs, retrieval-augmented frameworks have also been proposed for GFMs [18, 25, 47]. These methods typically retrieve subgraphs based on embeddings, conflating semantic proximity with topological relevance. In practice, nodes that are structurally interdependent may share little semantic resemblance, while semantically similar nodes are often topologically

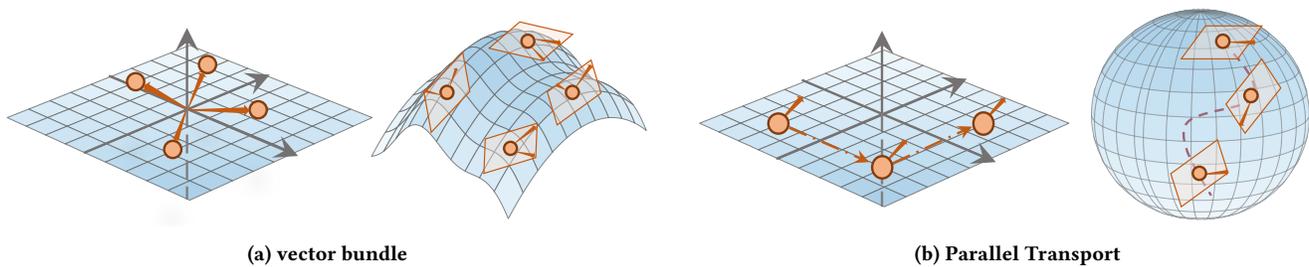


Figure 2: Illustration of parallel transport and vector bundle.

distant. Consequently, preserving structural dependencies becomes a critical bottleneck, yet it remains essential for achieving universal graph modeling. **Furthermore, the ability for graph generation remains conspicuously absent among state-of-the-art GFM (as shown in Table 1), despite generative modeling ability being a hallmark of LLMs.**

### 3 Riemannian Geometry Speaks Louder Than Words for Graphs

#### 3.1 Rethinking GFM

GFM represent a class of expressive, general-purpose, and pre-trainable neural architectures designed for graphs. At the heart of any foundation model lies the ability to identify generalities (or invariances) which enable broad adaptability. The success of LLMs is rooted in a shared word vocabulary that spans across texts. However, words alone fall short of capturing the structural complexity of graphs, as discussed above. The core challenge in the graph domain is thus to uncover the generalities that govern universal graph modeling and to establish a principled framework for understanding how structural semantics transfer across different graphs. **Standing at a crossroads, we propose to open a new direction for building GFM.** Before unfolding this blue-sky idea, we first return to the essence of graphs.

#### 3.2 New Geometric Toolkits

Graphs can be viewed as discrete analogs of manifolds, and Riemannian geometry provides a rigorous foundation for modeling them. This geometric framework naturally equips us with a powerful toolkit for GFM. We therefore argue that, **when it comes to graphs, Riemannian geometry speaks louder than words.**

**Manifolds and Vector Bundle.** A manifold is a space that locally resembles flat space but may curve globally, providing a natural geometric foundation for GFM to model the intrinsic structure of graphs beyond fixed coordinate assumptions. As shown in Figure 2a, a vector bundle assigns an independent vector space to each point on the manifold, allowing representations to reside in domain-specific vector spaces. This design enables flexible and expressive representation learning while preserving structural integrity, making it well-suited for multi-domain graph modeling.

#### **Riemannian Metric and Élie J. Cartan’s Moving Frame Method.**

A Riemannian metric equips a manifold with a smoothly varying inner product that defines local distances and angles, allowing the geometry to accommodate domain-specific structural patterns without distorting intrinsic graph structure. A moving frame provides a

local coordinate basis that evolves along the manifold and aligns representations by adapting to the local metric. Together, the metric and the moving frame allow the geometry to adapt across graph domains while maintaining consistent representations.

#### **Curvature, Levi-Civita Connection and Parallel Transports.**

A connection defines how vectors at different points on a manifold are compared. As shown in Figure 2b, parallel transport moves vectors between points while preserving geometric relationships, providing a principled geometric basis for message passing by transporting node representations into comparable spaces prior to aggregation. Furthermore, curvature measures the degree to which a manifold deviates from flatness. When representations are transported along different paths connecting the same nodes, the resulting curvature may differ, indicating geometric inconsistency in the underlying space. In this sense, curvature acts as an intrinsic geometric indicator of coherent inference.

#### 3.3 Blue Sky Idea

**Our Recommendation.** As introduced above, Riemannian geometry provides a rigorous foundation for graph modeling, illuminating a new direction for building GFM, a Riemannian geometry approach. In this paper, we put forward a blue sky idea of **Riemannian Foundation Model (RFM)**. This perspective shifts the focus of graph learning from designing representations to discovering geometry. Moving beyond the mere switching of representation spaces, RFM emphasizes graph intrinsic geometry such as flatness, symmetry, connectivity, and anisotropy. The key is to recover the manifold underlying graphs so that generality and structural invariance emerge naturally from the geometry itself, rather than being imposed through handcrafted architectures. This broader geometric perspective allows models to capture structural patterns that vary across domains within a unified mathematical framework. It offers a principled framework to uncover graph generalities, interpret knowledge transfer across domains, and endow models with endogenous capacities for structural inference and generation.

**Connection to Related Approaches.** We discuss the inherent difference among RFM, geometric deep learning and non-Euclidean graph learning. **Geometric deep learning** [4] focuses on explicit geometric data such as point clouds or meshes, modeling distances, angles, and 3D coordinates in Euclidean space. **Non-Euclidean graph learning** extends graph neural networks to curved spaces, such as hyperbolic or spherical manifolds, and has proven effective in modeling hierarchical and complex relational structures [10, 29]. This line of work centers on representation learning, and recent

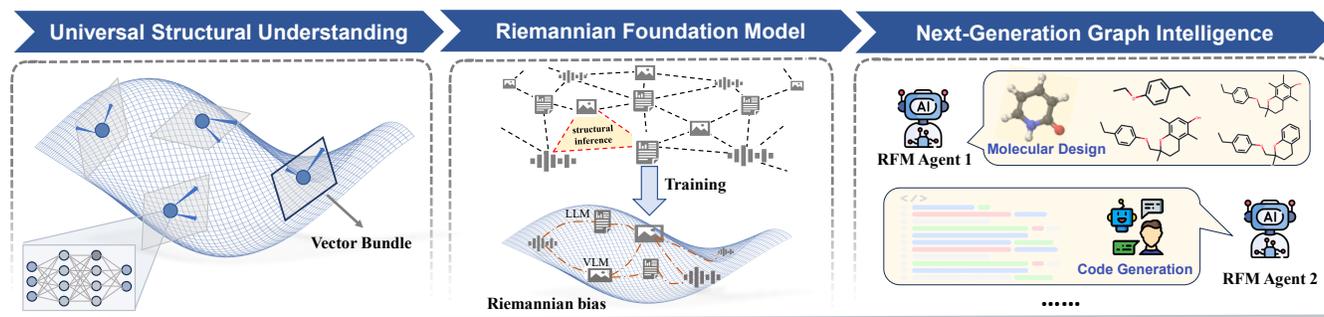


Figure 3: An illustration of the road to next-generation graph intelligence.

efforts introduce advanced manifold constructions tailored to specialized graphs [8, 26, 35]. In contrast, **RFM** emphasizes the intrinsic properties of the manifold itself such as flatness, symmetry and connectivity, independent of any embedding. Rather than operating in fixed coordinate spaces, RFM aims to uncover the underlying generalities and geometric structure governing real-world graphs across multiple domains.

## 4 A Progressive Agenda

Grounded in Riemannian geometry, we outline a progressive agenda towards universal graph modeling and beyond, as sketched in Figure 3. Specifically, we begin with universal structural understanding through intrinsic geometry, and then rebuild LLM with a Riemannian engine featured by the endogenous capability of structural inference, unlocking the next-generation graph intelligence.

### 4.1 Universal Structural Understanding

Here, we active a previously ignored concept of intrinsic geometry for pretraining and adaptation across the graph domains.

**Intrinsic Geometry & Invariant Structures.** We first address a fundamental objective in universal graph modeling: identifying generalizable patterns across graph domains. While extensive research has explored potential candidates for such generalities, including motifs [35], tree message-passing schemes [48, 49], graphons [58], and structural vocabularies [2, 18], these approaches often struggle with adaptability. Specifically, they lack a mechanism to determine which pretrained structural patterns remain applicable to specific target graphs. This often results in inconsistent functional behavior and suboptimal performance across different contexts. In contrast, Riemannian geometry provides the framework of vector bundles to capture intrinsic geometry, effectively uncovering structural invariants that persist across disparate domains.

**Principled Knowledge Transfer.** Second, we establish a principled framework to understand how structural semantics transfer across different graphs. Rather than relying on fixed geometries like hyperbolic or spherical spaces, RFM learns the smooth, underlying manifold directly from node representations to ensure a well-defined Levi-Civita connection. Given such a manifold, the concept of parallel transport can be leveraged to understand how knowledge is integrated or transferred across domains. Within this vector bundle construction, the derivation of invariant structures offers a measure of transfer effort and specifies the mechanism

by which behaviors learned from substructures during pretraining become applicable to unseen graphs.

### 4.2 Riemannian Foundation Model

We outline the blueprint of RFM for the future graph intelligence.

**Integrating LLMs.** We propose a unified path for integrating GFMs and LLMs into a cohesive whole. While universal graph modeling demands a deep understanding of both structural patterns and attribute semantics, these components exhibit distinct geometric properties: graph structures are inherently Riemannian, whereas attributes like text and images typically reside in Euclidean spaces. Simultaneously modeling such heterogeneous modalities remains a non-trivial task. To address this, we leverage the concept of vector bundles in Riemannian geometry. This principled framework reconciles diverse geometries by allowing LLMs to operate effectively within Euclidean fibers. Such a construction demonstrates that vector bundles, while preserving structural invariance, are also fundamentally compatible with the integration of LLMs.

**Endogenous Structural Inference.** We introduce endogenous structural inference as a hallmark of RFM. In our construction, large-scale post-training is conducted on a multimodal graph in the real world. By inherently anchoring linguistic tokens to a topological manifold, the generative process is thus aware of complex relationships among objects, equipped with the endogenous capability of structural inference, understanding the rationale underlying the linkages. While existing GFMs prioritize task-independent representations to achieve general-purpose utility, they often lack insight into the fundamental mechanisms of structure formation. By contrast, endogenous structural inference empowers RFM with enhanced ability in complex logic and structural reasoning.

**Graph Generation.** Here, we discuss graph generation within the framework of RFM, which remains conspicuously absent among current GFMs yet stands as a hallmark of LLMs. Graph generation is vital in a wide range of real-world applications ranging from drug discovery to circuit design. As claimed in the universality, the ability for Graph generation should be included in GFMs, yet it is evidently a distant goal. RFMs inherit the generative modeling ability during the integration of LLMs, and endogenous structural inference enables the model to recover the linguistic tokens as graphs. It not only achieves graph generation but also provides the interpretative clues for such generation.

**Our Message.** *The geometric toolkit provided by Riemannian geometry is ready to build GFM and redefine what a foundation model for graphs can be. RFM identifies the generalities using intrinsic geometry, provides principled interpretation of knowledge transfer, embodies the endogenous capability of structural inference and enables the graph generation which remains conspicuously absent among GFM. Also, we suggest a wider recognition of **intrinsic geometry** that governs the structural patterns in the graph domain, instead of the coordinates in the representation space, and emphasize that the **endogenous structural inference** should be a key to universal graph modeling and beyond.*

### 4.3 Graph Intelligence with RFM

Benefit from RFM, the evolution of AI demands a fundamental shift from linguistic pattern matching to intrinsic structural cognition, no longer viewing the world as mere isolated embeddings.

**Multimodal World Cognition beyond Graphs.** In our envisioned paradigm, the RFM serves as the fundamental medium across diverse modalities, driving next-generation AI from passive perception toward interactive spatial intelligence. As highlighted by recent World models [9, 12], true world cognition requires understanding not just what entities are, but how they physically and logically interact within a space. To achieve this, the RFM constructs a universal coordinate system from the environment’s underlying relational skeleton. Rather than loosely associating text, vision, and physical states through statistical similarity in an embedding space, the RFM anchors all disparate modalities directly onto this topological skeleton. Whether parsing the 3D spatial configuration of a robotic workspace or the logical prediction, the model enforces strict relational constraints. This structural grounding ensures that AI perception fundamentally aligns with human cognition.

**RFM Agents for Graph-structured Applications.** The application paradigm of graph learning will undergo a radical transition mirroring the evolution of LLMs: shifting from training isolated predictive models to deploying general-purpose autonomous agents [54]. In the next generation, the RFM functions as a central cognitive engine that interacts with the network environment and various soft tools. By perceiving the graph as an actionable space, the structural agent can navigate complex topologies, execute multi-step reasoning, and dynamically invoke structural tools to solve open-ended problems. For instance, in molecular discovery, rather than merely predicting a compound’s toxicity, an RFM agent can autonomously traverse the molecular graph to pinpoint toxic subgraphs and orchestrate chemical editing tools to swap functional groups. This transforms graph AI from a static analytical calculator into an interactive and decision-making entity capable of actively optimizing the structures it perceives.

## 5 Alternative Views

This section presents divergent perspective surrounding Graph Foundation Models. Current discussions primarily revolve around two central questions: the optimal architecture for graph intelligence, and the nature of its symbiosis with Large Language Models.

### 5.1 Architecture of Graph Foundation Model

Beyond message-passing GNNs, graph transformers have emerged as another prominent architecture for graph modeling, integrating structural inductive biases into standard self-attention mechanisms [15, 31, 41]. Recent efforts strive for universal graph modeling through the transformer architecture, exemplified by models such as GDL4LLM [67] and GraphFM [23]. These approaches conceptualize graphs as a new structural language, discretizing topological data into token sequences compatible with transformer backbones. Theoretically, such models can learn to simulate graph connectivity via massive-scale pre-training, suggesting that structural intelligence may emerge from compressing multi-domain features into a shared latent space. Meanwhile, other researchers attempt to combine text encoding with graph aggregation by nesting GNN components within transformer blocks [7, 30, 40]. From this perspective, the success of GFM hinges on sampling techniques and the efficiency of structural tokenization. However, we contend that this paradigm introduces a fundamental representational bottleneck: by forcing non-Euclidean graphs into a one-dimensional format, these models inevitably fragment global topology. This suggests that achieving a powerful GFM remains a distant goal and underscores the necessity of studying graph structures from a more native geometric perspective rather than treating them as sequences.

### 5.2 Collaboration with Large Language Model

Another debate centers on whether GFM should function as an external plugin for language models or serve as the intrinsic cognitive foundation for all reasoning. The prevailing paradigm maintains that LLMs should remain the primary engine of cognition. In this line, models either translate graph data into text-based prompts (e.g., GraphTranslator [62]) or project graph embeddings into the language semantic space (e.g., LLaGA [7]). While this offers flexibility, it relegates structural dependencies to secondary evidence and force them to be assimilated into the probabilistic patterns of human language, which inevitably creates a semantic bottleneck. In contrast, we argue that rather than serving as an external plugin, the impressive path forward lies in endogenous integration, unifying GFM and LLM into a cohesive whole. By inherently anchoring linguistic tokens to a topological manifold, the generative process becomes strictly bounded by structural dependencies grounded in the cognitive substrate of GFM. Such integration would empower next-generation large models to achieve enhanced capabilities for complex logic and reasoning.

## 6 Conclusion

The road to Graph Foundation Models is currently divided, with both prevailing approaches facing inherent challenges. This paper illuminates a new pathway—Riemannian geometry as the foundational principle for GFM. It is distinguished from current Riemannian approaches that mostly focus on the exogenous representation space orthogonal to cross-domain graph generalities. We introduce the blue sky idea of RFM featured by intrinsic geometry and endogenous structural inference. We showcase a progressive research agenda, redefining what a foundation model for graphs can be. Rather than incremental evolution, RFM heralds a transformative paradigm shift that leads the graph domain into its large model era.

## References

- [1] Beatrice Bevilacqua, Joshua Robinson, Jure Leskovec, and Bruno Ribeiro. 2025. Holographic Node Representations: Pre-training Task-Agnostic Node Embeddings. In *ICLR*.
- [2] Jianyuan Bo, Hao Wu, and Yuan Fang. 2025. Quantizing Text-attributed Graphs for Semantic-Structural Integration. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2*. 107–118.
- [3] Filip Bosnić and Mile Šikić. 2023. Finding Hamiltonian cycles with graph neural networks. In *2023 International Symposium on Image and Signal Processing and Analysis (ISPA)*. IEEE, 1–6.
- [4] Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. 2017. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine* 34, 4 (2017), 18–42.
- [5] Haibo Chen, Xin Wang, Zeyang Zhang, Haoyang Li, Ling Feng, and Wenwu Zhu. 2025. Autofgm: Automated graph foundation model with adaptive architecture customization. In *Forty-second International Conference on Machine Learning*.
- [6] Ming Chen, Zhewei Wei, Zengfeng Huang, Bolin Ding, and Yaliang Li. 2020. Simple and deep graph convolutional networks. In *International conference on machine learning*. PMLR, 1725–1735.
- [7] Runjin Chen, Tong Zhao, Ajay Kumar Jaiswal, Neil Shah, and Zhangyang Wang. 2024. LLaGA: Large Language and Graph Assistant. In *International Conference on Machine Learning*. PMLR, 7809–7823.
- [8] Cheng Deng, Fan Xu, Jiaying Ding, Luoyi Fu, Weinan Zhang, and Xinbing Wang. 2023. FMGNN: Fused manifold graph neural network. *arXiv preprint arXiv:2304.01081* (2023).
- [9] Jingtao Ding, Yunke Zhang, Yu Shang, Yuheng Zhang, Zefang Zong, Jie Feng, Yuan Yuan, Hongyuan Su, Nian Li, Nicholas Sukiennik, et al. 2025. Understanding world or predicting future? a comprehensive survey of world models. *Comput. Surveys* 58, 3 (2025), 1–38.
- [10] Manfredo Perdigão Do Carmo and J Flaherty Francis. 1992. *Riemannian geometry*. Vol. 393. Springer.
- [11] Moshe Eliasof, Krishna Sri Ipsit Mantri, Beatrice Bevilacqua, Bruno Ribeiro, and Carola-Bibiane Schönlieb. [n. d.]. Bridging Input Feature Spaces Towards Graph Foundation Models. In *The Fourteenth International Conference on Learning Representations*.
- [12] David Ha and Jürgen Schmidhuber. 2018. World models. *arXiv preprint arXiv:1803.10122* 2, 3 (2018), 440.
- [13] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. *Advances in neural information processing systems* 30 (2017).
- [14] Yufei He and Bryan Hooi. 2024. UniGraph: Learning a Cross-Domain Graph Foundation Model From Natural Language. *arXiv* (2024).
- [15] Yufei He, Zhenyu Hou, Yukuo Cen, Jun Hu, Feng He, Xu Cheng, Jie Tang, and Bryan Hooi. 2025. Generalizing Graph Transformers Across Diverse Graphs and Tasks via Pre-Training. *IEEE Transactions on Knowledge and Data Engineering* 38, 2 (2025), 1114–1128.
- [16] Zhenyu Hou, Haozhan Li, Yukuo Cen, Jie Tang, and Yuxiao Dong. 2024. Graphalign: Pretraining one graph neural network on multiple graphs via feature alignment. *arXiv* (2024).
- [17] Yuntong Hu, Zheng Zhang, and Liang Zhao. [n. d.]. Beyond Text: A Deep Dive into Large Language Models' Ability on Understanding Graph Data. In *NeurIPS 2023 Workshop: New Frontiers in Graph Learning*.
- [18] Xinke Jiang, Rihong Qiu, Yongxin Xu, Yichen Zhu, Ruizhe Zhang, Yuchen Fang, Chu Xu, Junfeng Zhao, and Yasha Wang. 2024. Ragraph: A general retrieval-augmented graph learning framework. *Advances in Neural Information Processing Systems* 37 (2024), 29948–29985.
- [19] Nicolas Keriven. 2022. Not too little, not too much: a theoretical analysis of graph (over) smoothing. *NeurIPS* (2022).
- [20] Jinwoo Kim, Xingyue Huang, Krzysztof Olejniczak, Kyungbin Min, Michael Bronstein, Seunghoon Hong, and İsmail İlkan Ceylan. 2025. Flock: A Knowledge Graph Foundation Model via Learning on Random Walks. *arXiv preprint arXiv:2510.01510* (2025).
- [21] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
- [22] Kerstin Klaser, Blazej Banaszczyk, Samuel Maddrell-Mander, Callum McLean, Luis Müller, Ali Parviz, Shenyang Huang, and Andrew W Fitzgibbon. 2024. Minimol: A parameter-efficient foundation model for molecular learning. In *ICML 2024 Workshop*.
- [23] Divyansha Lachi, Mehdi Azabou, Vinam Arora, and Eva Dyer. 2024. GraphFM: A Scalable Framework for Multi-Graph Pretraining. *arXiv* (2024).
- [24] Chaojian Li, Zhifan Ye, {Massimiliano Lupò} Pasini, {Jong Youl} Choi, Cheng Wan, and Prasanna Balaprakash. 2025. In *2025 62nd ACM/IEEE Design Automation Conference, DAC 2025 (Proceedings - Design Automation Conference)*. Institute of Electrical and Electronics Engineers Inc. doi:10.1109/DAC63849.2025.11132864 Publisher Copyright: © 2025 IEEE.; 62nd ACM/IEEE Design Automation Conference, DAC 2025 ; Conference date: 22-06-2025 Through 25-06-2025.
- [25] Jintang Li, Ruofan Wu, Yuchang Zhu, Huizhe Zhang, Liang Chen, and Zibin Zheng. 2025. Are Large Language Models In-Context Graph Learners? *arXiv preprint arXiv:2502.13562* (2025).
- [26] Linyu Li, Zhi Jin, Xuan Zhang, Haoran Duan, Jishu Wang, Zhengwei Tao, Haiyan Zhao, and Xiaofeng Zhu. 2025. Multi-view riemannian manifolds fusion enhancement for knowledge graph completion. *IEEE Transactions on Knowledge and Data Engineering* (2025).
- [27] Hao Liu, Jiarui Feng, Lecheng Kong, Ningyue Liang, Dacheng Tao, Yixin Chen, and Muhan Zhang. 2024. One for all: Towards training one graph model for all classification tasks. In *ICLR*.
- [28] Jingzhe Liu, Haitao Mao, Zhikai Chen, Wenqi Fan, Mingxuan Ju, Tong Zhao, Neil Shah, and Jiliang Tang. 2024. One Model for One Graph: A New Perspective for Pretraining with Cross-domain Graphs. *arXiv* (2024).
- [29] Qi Liu, Maximilian Nickel, and Douwe Kiela. 2019. Hyperbolic graph neural networks. *Advances in neural information processing systems* 32 (2019).
- [30] Zheyuan Liu, Xiaoxin He, Yijun Tian, and Nitesh V Chawla. 2024. Can we soft prompt LLMs for graph learning tasks?. In *WWW*.
- [31] Ladislav Rampasek, Mikhail Galkin, Vijay Prakash Dwivedi, Anh Tuan Luu, Guy Wolf, and Dominique Beaini. 2022. Recipe for a General, Powerful, Scalable Graph Transformer. In *NeurIPS*.
- [32] Junhua Shi, Qingyun Sun, Haonan Yuan, and Xingcheng Fu. 2025. SA<sup>2</sup> GFM: Enhancing Robust Graph Foundation Models with Structure-Aware Semantic Augmentation. *arXiv preprint arXiv:2512.07857* (2025).
- [33] Li Sun, Zhenhao Huang, Silei Chen, Lanxu Yang, Junda Ye, Sen Su, and Philip S Yu. [n. d.]. Multi-Domain Transferable Graph Gluing for Building Graph Foundation Models. In *The Fourteenth International Conference on Learning Representations*.
- [34] Li Sun, Zhenhao Huang, Hao Peng, Yujie Wang, Chunyang Liu, and Philip S Yu. 2024. Lsenet: Lorentz structural entropy neural network for deep graph clustering. *arXiv preprint arXiv:2405.11801* (2024).
- [35] Li Sun, Zhenhao Huang, Suyang Zhou, Qiqi Wan, Hao Peng, and Philip Yu. 2025. Riemannfm: Learning a graph foundation model from riemannian geometry. In *Proceedings of the ACM on Web Conference 2025*. 1154–1165.
- [36] Li Sun, Zhongbao Zhang, Jiawei Zhang, Feiyang Wang, Hao Peng, Sen Su, and Philip S Yu. 2021. Hyperbolic variational graph neural network for modeling dynamic graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 4375–4383.
- [37] Yuanfu Sun, Zhengnan Ma, Yi Fang, Jing Ma, and Qiaoyu Tan. 2025. Graphicl: Unlocking graph learning potential in llms through structured prompt design. In *Findings of the Association for Computational Linguistics: NAACL 2025*. 2440–2459.
- [38] Yifei Sun, Yang Yang, Xiao Feng, Zijun Wang, Haoyang Zhong, Chunping Wang, and Lei Chen. 2025. Handling feature heterogeneity with learnable graph patches. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 1*. 1313–1324.
- [39] Maciej Sypetkowski, Frederik Wenkel, Farimah Poursafaei, Nia Dickson, Karush Suri, Philip Fradkin, and Dominique Beaini. 2025. On the scalability of gnns for molecular graphs. *NeurIPS* (2025).
- [40] Jiabin Tang, Yuhao Yang, Wei Wei, Lei Shi, Lixin Su, Suqi Cheng, Dawei Yin, and Chao Huang. 2024. Graphgpt: Graph instruction tuning for large language models. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 491–500.
- [41] Jiabin Tang, Yuhao Yang, Wei Wei, Lei Shi, Long Xia, Dawei Yin, and Chao Huang. 2024. Higt: Heterogeneous graph language model. In *KDD*.
- [42] Ziyuan Tang and Jie Chen. 2025. Toward a Graph Foundation Model: Pre-Training Transformers With Random Walks. *arXiv preprint arXiv:2506.14098* (2025).
- [43] Loring W Tu. 2011. Manifolds. In *An Introduction to Manifolds*. Springer, 47–83.
- [44] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. In *International Conference on Learning Representations*.
- [45] Feiyang Wang, Li Sun, and Zhongbao Zhang. 2020. Hyperbolic user identity linkage across social networks. In *GLOBECOM 2020-2020 IEEE global communications conference*. IEEE, 1–6.
- [46] Shuo Wang, Bokui Wang, Zhixiang Shen, Boyan Deng, and Zhao Kang. 2025. Multi-Domain Graph Foundation Models: Robust Knowledge Transfer via Topology Alignment. In *International Conference on Machine Learning*. PMLR, 64806–64821.
- [47] Xingliang Wang, Zemin Liu, Junxiao Han, and Shuiguang Deng. 2025. RAG4GFM: Bridging Knowledge Gaps in Graph Foundation Models through Graph Retrieval Augmented Generation. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- [48] Zehong Wang, Zheyuan Zhang, Nitesh V Chawla, Chuxu Zhang, and Yanfang Ye. 2024. GFT: Graph Foundation Model with Transferable Tree Vocabulary. In *NeurIPS*.
- [49] Zehong Wang, Zheyuan Zhang, Tianyi Ma, Nitesh V Chawla, Chuxu Zhang, and Yanfang Ye. 2025. Towards Graph Foundation Models: Learning Generalities Across Graphs via Task-Trees. In *Proceedings of the Forty-Second International Conference on Machine Learning*.

- [50] Zhengyu Wu, Yinlin Zhu, Xunkai Li, Ziang Qiu, Rong-Hua Li, Guoren Wang, and Chenghu Zhou. 2025. FedBook: A Unified Federated Graph Foundation Codebook with Intra-domain and Inter-domain Knowledge Modeling. *arXiv preprint arXiv:2510.07755* (2025).
- [51] Lianghao Xia and Chao Huang. 2024. Anygraph: Graph foundation model in the wild. *arXiv preprint arXiv:2408.10700* (2024).
- [52] Lianghao Xia, Ben Kao, and Chao Huang. 2024. Opengraph: Towards open graph foundation models. In *EMNLP*.
- [53] Yuhao Xu, Xinqi Liu, Keyu Duan, Yi Fang, Yu-Neng Chuang, Daochen Zha, and Qiaoyu Tan. 2024. GraphFM: A Comprehensive Benchmark for Graph Foundation Model. *arXiv preprint arXiv:2406.08310* (2024).
- [54] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. 2022. React: Synergizing reasoning and acting in language models. In *The eleventh international conference on learning representations*.
- [55] Ruosong Ye, Caiqi Zhang, Runhui Wang, Shuyuan Xu, and Yongfeng Zhang. 2024. Language is all a graph needs. In *Findings of the association for computational linguistics: EACL 2024*. 1955–1973.
- [56] Xingtong Yu, Zechuan Gong, Chang Zhou, Yuan Fang, and Hui Zhang. 2025. Samgpt: Text-free graph foundation model for multi-domain pre-training and cross-domain adaptation. In *Proceedings of the ACM on Web Conference 2025*. 1142–1153.
- [57] Xingtong Yu, Chang Zhou, Yuan Fang, and Xinming Zhang. 2024. Text-free multi-domain graph pre-training: Toward graph foundation models. *arXiv preprint arXiv:2405.13934* (2024).
- [58] Haonan Yuan, Qingyun Sun, Junhua Shi, Xingcheng Fu, Bryan Hooi, Jianxin Li, and Philip S Yu. 2025. GRAVER: Generative Graph Vocabularies for Robust Graph Foundation Models Fine-tuning. *arXiv preprint arXiv:2511.05592* (2025).
- [59] Haonan Yuan, Qingyun Sun, Junhua Shi, Xingcheng Fu, Bryan Hooi, Jianxin Li, and Philip S Yu. 2025. How much can transfer? BRIDGE: Bounded multi-domain graph foundation model with generalization guarantees. In *Forty-second International Conference on Machine Learning*.
- [60] Jiawei Zhang. 2023. Graph-toolformer: To empower llms with graph reasoning ability via prompt augmented by chatgpt. *arXiv preprint arXiv:2304.11116* (2023).
- [61] Juzheng Zhang, Yatao Bian, Yongqiang Chen, and Quanming Yao. 2024. Unimot: Unified molecule-text language model with discrete token representation. *arXiv* (2024).
- [62] Mengmei Zhang, Mingwei Sun, Peng Wang, Shen Fan, Yanhu Mo, Xiaoxiao Xu, Hong Liu, Cheng Yang, and Chuan Shi. 2024. Graphtranslator: Aligning graph model to large language model for open-ended tasks. In *Proceedings of the ACM Web Conference 2024*. 1003–1014.
- [63] Haihong Zhao, Aochuan Chen, Xiangguo Sun, Hong Cheng, and Jia Li. 2024. All in one and one for all: A simple yet effective method towards cross-domain graph pretraining. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 4443–4454.
- [64] Jianan Zhao, Zhaocheng Zhu, Mikhail Galkin, Hesham Mostafa, Michael M Bronstein, and Jian Tang. 2025. Fully-inductive Node Classification on Arbitrary Graphs. In *ICLR*.
- [65] Jianan Zhao, Le Zhuo, Yikang Shen, Meng Qu, Kai Liu, Michael Bronstein, Zhaocheng Zhu, and Jian Tang. 2023. Graphtext: Graph reasoning in text space. *arXiv preprint arXiv:2310.01089* (2023).
- [66] Zihao Zhao, Xinlong Zhai, Jinyu Yang, and Chuan Shi. 2025. Towards Text-free Graph Foundation Models: Rethinking Multi-Domain Graph Contrastive Learning. *arXiv preprint arXiv:2506.22510* (2025).
- [67] Huachi Zhou, Jiahe Du, Chuang Zhou, Chang Yang, Yilin Xiao, Yuxuan Xie, and Xiao Huang. 2025. Each Graph is a New Language: Graph Learning with LLMs. *arXiv* (2025).
- [68] Yinlin Zhu, Xunkai Li, Jishuo Jia, Miao Hu, Di Wu, and Meikang Qiu. 2025. Towards effective federated graph foundation model via mitigating knowledge entanglement. *arXiv preprint arXiv:2505.12684* (2025).
- [69] Yun Zhu, Haizhou Shi, Xiaotang Wang, Yongchao Liu, Yaoke Wang, Boci Peng, Chuntao Hong, and Siliang Tang. 2025. Graphclip: Enhancing transferability in graph foundation models for text-attributed graphs. In *Proceedings of the ACM on Web Conference 2025*. 2183–2197.